arXiv:1612.05478v2 [cs.CV] 5 Jan 2017

Video Propagation Networks

Varun Jampani[1], Raghudeep Gadde[1,2] and Peter V. Gehler[1,2]
[1]Max Planck Institute for Intelligent Systems, Tübingen, Germany
[2]Bernstein Center for Computational Neuroscience, Tübingen, Germany
{varun.jampani,raghudeep.gadde,peter.gehler}@tuebingen.mpg.de

Abstract

In this paper we propose a technique that propagates information forward through video data. The method is conceptually simple and can be applied to tasks that require the propagation of structured information, such as semantic labels, based on video content. We propose a Video Propagation Network *that processes video frames in an adaptive manner. The model is applied online: it propagates information forward without the need to access future frames other than the current ones. In particular we combine two components, a temporal bilateral network for dense and video adaptive filtering, followed by a spatial network to refine features and increased flexibility. We present experiments on video object segmentation and semantic video segmentation and show increased performance comparing to the best previous task-specific methods, while having favorable runtime. Additionally we demonstrate our approach on an example regression task of propagating color in a grayscale video.*

1. Introduction

In this work, we focus on the problem of propagating structured information across video frames. This problem appears in many forms (e.g., semantic segmentation or depth estimation) and is a pre-requisite for many applications. An example instance is shown in Fig. 1. Given an accurate object mask for the first frame, the problem is to propagate this mask forward through the entire video sequence. Propagation of semantic information through time and video colorization are other problem instances.

Videos pose both technical and representational challenges. The presence of scene and camera motion lead to the difficult association problem of optical flow. Video data is computationally more demanding than static images. A naive per-frame approach would scale at least linear with frames. These challenges complicate the use of standard convolutional neural networks (CNNs) for video processing. As a result, many previous works for video propagation use slow optimization based techniques.

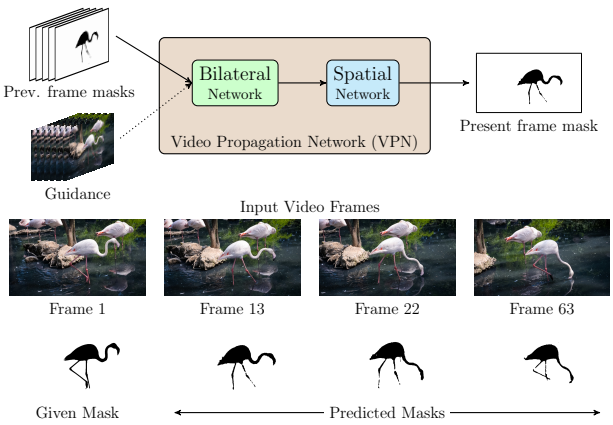

Figure 1. **Video Propagation with VPNs.** The end-to-end trained VPN network is composed of a bilateral network followed by a standard spatial network and can be used for propagating information across frames. Shown here is an example result of foreground mask from the 1^{st} frame to other video frames.

We propose a generic neural network architecture that propagates information across video frames. The main innovation is the use of image adaptive convolutional operations that automatically adapt to the video stream content. This allows the network to adapt to the changing content of the video stream. It can be applied to several types of information, e.g., labels, colors, etc. and runs online, that is, only requiring current and previous frames.

Our architecture is composed of two components (see Fig. 1). A temporal *bilateral network* that performs image-adaptive spatial-temporal dense filtering. This part allows to connect densely all pixels from current and previous frames and to propagate associated pixel information to the current frame. The bilateral network allows the specification of a metric between video pixels and allows a straight-forward integration of temporal information. This is followed by a standard *spatial CNN* on the filter output to refine and predict for the present video frame. We call this combination a *Video Propagation Network (VPN)*. In effect we are combining a filtering technique with rather small spatial CNNs

which leads to a favorable runtime compared to many previous approaches.

VPNs have the following suitable properties for video processing:

General applicability: VPNs can be used for propagating any type of information content i.e., both discrete (e.g., semantic labels) and continuous (e.g., color) information across video frames.

Online propagation: The method needs no future frames and so can be used for online video analysis.

Long-range and image adaptive: VPNs can efficiently handle a large number of input frames and are adaptive to the video.

End-to-end trainable: VPNs can be trained end-to-end, so they can be used in other deep network architectures.

Favorable runtime: VPNs have favorable runtime in comparison to several current best methods, also making them amenable for learning with large datasets.

Empirically we show that VPNs, despite being generic, perform better than published approaches on video object segmentation and semantic label propagation while being faster. VPNs can easily be integrated into sequential per-frame approaches and require only a small fine-tuning step that can be performed separately.

2. Related Work

The literature on propagating information across video frames contains a vast and varied number of approaches. Here, we only discuss those works that are related to our technique and applications.

General propagation techniques Techniques for propagating content across image/video pixels are predominantly optimization based or filtering techniques. Optimization based techniques typically formulate the propagation as an energy minimization problem on a graph constructed across video pixels or frames. A classic example is the color propagation technique from [46] which uses graph structure that encodes prior knowledge about pixel colors in a local neighborhood. Although efficient closed-form solutions [47] exists for certain scenarios, optimization tends to be slow due to either large graph structures for videos and/or the use of complex connectivity resulting in the use of iterative optimization schemes. Fully-connected conditional random fields (CRFs) [41] open a way for incorporating dense and long-range pixel connections while retaining fast inference.

Filtering techniques [40, 15, 30] aim to propagate information with the use of image or video filters resulting in fast runtimes compared to optimization techniques. Bilateral filtering [5, 73] is one of the popular filters for long-range information propagation. A popular application is joint bilateral upsampling [40] that upsamples a low-resolution signal with the use of a high-resolution guidance image. The works of [51, 22, 37, 34, 82, 66] showed that one can back-propagate through the bilateral filtering operation for learning filter parameters [37, 34] or doing optimization in the bilateral space [8, 7]. Recently, several works proposed to do upsampling in images by learning CNNs that mimic edge-aware filtering [79] or that directly learn to upsample [49, 32]. Most of these works are confined to images and are either not extendable or computationally too expensive for videos. We leverage some of these previous works and propose a scalable yet robust neural network based approach for video content propagation. We will discuss more about bilateral filtering, that forms the core of our approach, in Section 3.

Video object segmentation Prior work on video object segmentation can be broadly categorized into two types: Semi-supervised methods that require manual annotation to define what is foreground object and unsupervised methods that does segmentation completely automatically. Unsupervised techniques such as [25, 48, 45, 55, 78, 81, 72, 23] use some prior information about the foreground objects such as distinctive motion, saliency etc. And, they typically fail if these assumptions do not hold in a video.

In this work, we focus on the semi-supervised task of propagating the foreground mask from the first frame to the entire video. Existing works predominantly use graph-based optimization frameworks that perform graph-cuts [9, 10, 69] on video data. Several of these works [64, 50, 61, 77, 39, 33] aim to reduce the complexity of graph structure with clustering techniques such as spatial-temporal superpixels and optical flow [76]. Another direction was to estimate correspondence between different frame pixels [4, 6, 44] by using nearest neighbor fields [26] or optical flow [18] and then refine the propagated masks with the use of local classifiers. Closest to our technique are the works of [60] and [53]. [60] proposed to use fully-connected CRF over the refined object proposals across the video frames. [53] proposed a graph-cut in the bilateral space. Our approach is similar in the regard that we also use a bilateral space embedding. Instead of graph-cuts, we learn propagation filters in the high-dimensional bilateral space with CNNs. This results in a more generic architecture and allows integration into other deep learning frameworks.

Two contemporary works [14, 36] proposed CNN based approaches for video object segmentation. Both works rely on fine-tuning a deep network using the first frame annotation of a given test sequence. This could potentially result in overfitting to the test background. In contrast, the proposed approach relies only on offline training and thus can be easily adapted to different problem scenarios.

Semantic video segmentation Earlier methods such as [12, 70] use structure from motion on video frames to compute geometrical and/or motion features. More recent works [24, 16, 19, 54, 74, 43] construct large graph-

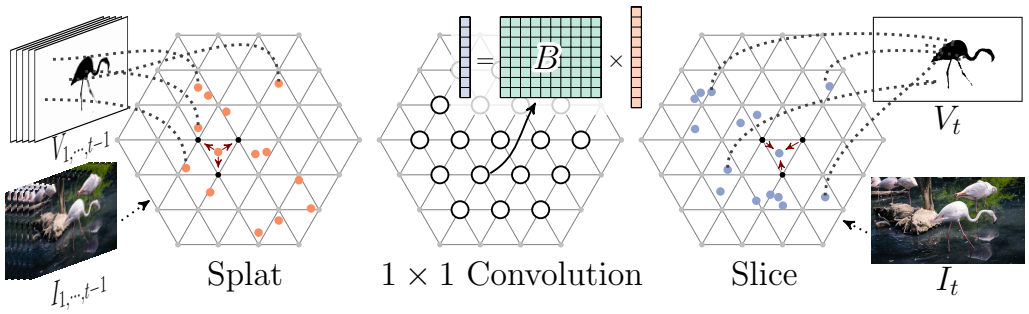

Figure 2. **Schematic of Fast Bilateral Filtering for Video Processing.** Mask probabilities from previous frames $V_{1,\cdots,t-1}$ are splatted on to the lattice positions defined by the image features $f_{I_1}, f_{I_2}, \cdots, f_{I_{t-1}}$. The splatted result is convolved with a 1×1 filter B, and the filtered result is sliced back to the original image space to get V_t for the present frame. Input and output need not be V_t, but can also be an intermediate neural network representation. B is learned via back-propagation through these operations.

ical models on videos and enforce temporal consistency across frames. [16] used dynamic temporal links in their CRF energy formulation. [19] proposes to use Perturb-and-MAP random field model with spatial-temporal energy terms based on Potts model and [54] propagate predictions across time by learning a similarity function between pixels of consecutive frames.

In the recent years, there is a big leap in the performance of semantic image segmentation [52, 17] with the use of CNNs but mostly applied to images. Recently, [67] proposed to retain the intermediate CNN representations while sliding the image based CNN across the frames. Another approach, which inspired our work, is to take unary predictions from CNN and then propagate semantic information across the frames. A recent prominent approach in this direction is of [43] which proposes a technique for optimizing feature spaces for fully-connected CRF.

3. Bilateral Filtering

We briefly review the background of bilateral filtering and the extensions that we will need to build a Bilateral Network for videos. Bilateral filtering has its roots in image denoising [5, 73] and has been developed as an edge-preserving filter. It has found numerous applications [58] and recently found its way into neural network architectures [82, 27]. In this work, we will use this filtering at the core of our approach and make use of the image/video-adaptive connectivity as a way to cope with scenes in motion. The most used bilateral filter kernel is a Gaussian filter, we make use of the recent work of [34] that generalizes this filter to a full parameterized operation.

Bilateral filtering of a vector V can be viewed as a matrix-vector multiplication operation with a matrix W

$$\hat{V}(i) = \sum_j W(i,j)V(j), \tag{1}$$

where the filter weights $W(i,j)$ depend on features $\mathbf{f}_i, \mathbf{f}_j$ of every input index i, j. For example a Gaussian bilateral filter amounts to a particular choice of W as $W(i,j) = \frac{1}{Z}\exp\left(-\frac{1}{2}(\mathbf{f}_i - \mathbf{f}_j)^\top \Sigma^{-1}(\mathbf{f}_i - \mathbf{f}_j)\right)$. The choice of features $\mathbf{f}$ define the effect of the filter, the way it adapts to image content. To use only positional features, $\mathbf{f} = (x, y)$, the bilateral filter operation reduces to a spatial Gaussian filter, with width controlled by Σ. A common choice for edge-preserving filtering is to choose color and position features $\mathbf{f} = (x, y, r, g, b)^\top$. This results in image smoothing without blurring across the edges.

The filter values W change for every pixel pairs i, j and depend on the image/video content. Therefore, a naive implementation of the vector-product computation in Eq. 1 is prohibitive. Due to the importance of this filtering operation, several fast algorithms [2, 3, 57, 28] have been proposed, that make use of the structure of W and V, not building up the matrix altogether but computing the multiplication result directly. One natural view that inspired several implementations was offered by [57], who viewed the bilateral filtering operation as a computation in a higher dimensional space. Their observation was that filtering can be implemented by 1. projecting V into a high-dimensional grid (*splatting*), 2. high-dimensional filtering (*convolving*) and 3. projecting down the result at the points of interest (*slicing*). The high-dimensional grid is also called *bilateral space/grid*. All these three operations are linear

$$\hat{V} = S_{slice} B S_{splat} V, \tag{2}$$

where, S_{splat} and S_{slice} denotes the mapping to-from image pixels and bilateral grid, and B denotes convolution (traditionally Gaussian) in the bilateral space. The problem with this approach is that a standard $n \times n \times n \times \ldots$ convolution on a regular grid requires handling of an exponential number of grid points. This was circumvented by a special data structure, the permutohedral lattice as proposed in [2]. Effectively permutohedral filtering scales linearly with dimension, resulting in fast execution time.

The recent work of [37, 34] then generalized the bilateral filter in the permutohedral lattice and demonstrated how it

can be learned via back-propagation. This allowed the construction of image-adaptive filtering operations into deep learning architectures, which we will build upon. See Fig. 2 for a 2D illustration of permutohedral lattices. Refer to [2] for more details on bilateral filtering using permutohedral lattice and refer to [34] for details on learning general permutohedral filters via back-propagation.

4. Video Propagation Networks

We aim to adapt the bilateral filtering operation to predict information forward in time, across video frames. Formally, we work on a sequence of n (color or grayscale) images $\mathbf{I} = \{I_1, I_2, \cdots, I_n\}$ and denote with $\mathbf{V} = \{V_1, V_2, \cdots, V_n\}$ a sequence of outputs, one per frame. Consider as an example a sequence $V_1, \ldots, V_n$ of foreground masks for a moving object in the scene. Our goal is to develop an online propagation method, that is, a function that has no access to the future frames. Formally we predict V_t, having observed the video up to frame t and possibly previous $V_{1,\cdots,t-1}$

$$\mathcal{F}(V_{t-1}, V_{t-2}, \cdots; I_t, I_{t-1}, I_{t-2}, \cdots) = V_t. \quad (3)$$

If training examples (I, V) with full or partial knowledge of V are available, it is possible to learn $\mathcal{F}$ and for a complex and unknown relationship between input and output, a deep CNN is a natural design choice. However, any learning based method has to face the main challenge: the scene and camera motion and its effect on V. Since no motion in two different videos is the same, fixed-sized static receptive fields of CNN units are insufficient. We propose to resolve this with video-adaptive convolutional component, an adaption of the bilateral filtering to videos. Our Bilateral Network (Section 4.1) has a connectivity that adapts to video sequences, its output is then fed into a common Spatial Network (Section 4.2) that further refines the desired output. The combined network layout of this Video Propagation Network is depicted in Fig. 3. It is a sequence of learnable bilateral and spatial filters that is efficient, trainable end-to-end and adaptive to the video input.

4.1. Bilateral Network (BNN)

In this section, we describe the extension of the learnable bilateral filtering to video data. Several properties of bilateral filtering make it a perfect candidate for information propagation in videos. In particular, our method is inspired by two main ideas that we extend in this work: joint bilateral upsampling [40] and learnable bilateral filters [34]. Although, bilateral filtering has been used for filtering video data before [56], its use has been limited to fixed filter weights (say, Gaussian).

Fast Bilateral Upsampling across Frames The idea of joint bilateral upsampling [40] is to view upsampling as a filtering operation. A high resolution guidance image is used to up-sample a low-resolution result. In short, a smaller number of input points are given $V_i, \mathbf{f}_i; i = 1, \ldots, N_{in}$, for example a segmentation result V_i at a lower resolution. This is then scaled to a larger number of output points $\mathbf{f}_j; j = 1, \ldots, N_{out}$ using the bilateral filtering operation, that is to compute Eq. 1, where the sum runs over all N_{in} points and the output is computed for all N_{out} positions. We will use this idea to propagate content from previous frames to the current frame (all of which have the same dimensions), using the current frame as a guidance image. This is illustrated in Fig. 2. We take all previous frame results $V_{1,\cdots,t-1}$ and splat them into a lattice using the features computed on video frames $I_{1,\cdots,t-1}$. A filtering (described below) is applied to every lattice point and the result is then sliced back using the current frame I_t. This result need not be the final V_t, in fact we compute a filter bank of responses and continue with further processing as will be discussed.

For videos, we need to extend bilateral filtering to temporal data, and there are two natural choices. First, one can simply attach a frame index t as an additional time dimension to the input data, yielding a six dimensional feature vector $\mathbf{f} = (x, y, r, g, b, t)^\top$ for every pixel in every frame. The summation in Eq. 1 now runs over *all* previous frames and pixels. Imagine a video where an object moves to reveal some background. Pixels of the object and background will be close spatially (x, y) and temporally (t) but likely be of different color (r, g, b). Therefore they will have no strong influence on each other (being splatted to distant positions in the six-dimensional bilateral space). In summary, one can understand the filter to be adaptive to color changes across frames, only pixels that are static and have similar color have a strong influence on each other (end up nearby in the lattice space). The second possibility is to use optical flow. If the perfect flow is available, the video frames could be warped into a common frame of reference. This would resolve the corresponding problem and make information propagation much easier. We can make use of an optical flow estimate by warping pixel positions (x, y) by their displacement vector (u_x, u_y) to $(x + u_x, y + u_y)$.

Another property of permutohedral filtering that we exploit is that the *input points need not lie on a regular grid* since the filtering is done in the high-dimensional lattice. Instead of splatting millions of pixels on to the lattice, we randomly sample or use superpixels and perform filtering using these sampled points as input to the filter. In practice, we observe that this results in big computational gains with minor drop in performance (more in Section 5.1).

Learnable Bilateral Filters The property of propagating information forward using a guidance image through filtering solves the problem of pixel association. But a Gaussian filter may be insufficient and further, we would like to increase the capacity by using a filter bank instead of a single fixed filter. We propose to use the technique of [34]

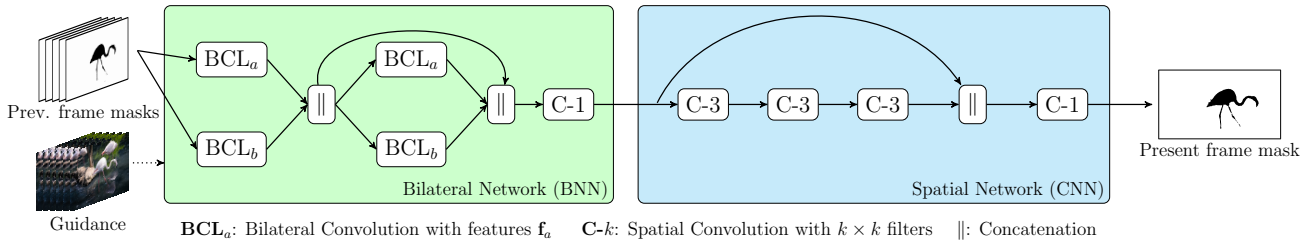

Figure 3. **Computation Flow of Video Propagation Network.** Bilateral networks (BNN) consist of a series of bilateral filterings interleaved with ReLU non-linearities. The filtered information from BNN is then passed into a spatial network (CNN) which refines the features with convolution layers interleaved with ReLU non-linearities, resulting in the prediction for the current frame.

to learn the filter values in the permutohedral lattice using back-propagation.

The process works as follows. A input video is used to determine the positions in the bilateral space to splat the input points $V(i) \in \mathbb{R}^D$ i.e., the features $\mathbf{f}$ (e.g. (x, y, r, g, b, t)) define the splatting matrix S_{splat}. This leads to a number of vectors $V_{splatted} = S_{splat}V$, that lie on the permutohedral lattice, with dimensionality $V_{splatted} \in \mathbb{R}^D$. In effect, the splatting operation groups points that are close together, that is, they have similar $\mathbf{f}_i, \mathbf{f}_j$. All lattice points are now filtered using a filter bank $B \in \mathbb{R}^{F \times D}$ which results in F dimensional vectors on the lattice points. These are sliced back to the N_{out} points of interest (present video frame). The values of B are learned by back-propagation. General parametrization of B from [34, 37] allows to have any neighborhood size for the filters. Since constructing the neighborhood structure in high-dimensions is time consuming, we choose to use 1×1 filters for speed reasons. This makes up one *Bilateral Convolution Layer (BCL)* which we will stack and concatenate to form a Bilateral Network. See Fig. 2 for an illustration of a BCL.

BNN Architecture The Bilateral Network (BNN) is illustrated in the green box of Fig. 3. The input is a video sequence $\mathbf{I}$ and the corresponding predictions $\mathbf{V}$ up to frame t. Those are filtered using two BCLs with 32 filters each. For both BCLs, we use the same features $\mathbf{f}$ but scale them with different diagonal matrices $\mathbf{f}_a = \Lambda_a\mathbf{f}, \mathbf{f}_b = \Lambda_b\mathbf{f}$. The feature scales are found by cross-validation. The two 32 dimensional outputs are concatenated, passed through a ReLU non-linearity and passed to a second layer of two separate BCL filters that uses same feature spaces $\mathbf{f}_a, \mathbf{f}_b$. The output of the second filter bank is then reduced using a 1×1 spatial filter (C-1) to map to the original dimension of V. We investigated scaling frame inputs with an exponential time decay and found that, when processing frame t, a re-weighting with $(\alpha V_{t-1}, \alpha^2 V_{t-2}, \alpha^3 V_{t-3} \cdots)$ with $0 \leq \alpha \leq 1$ improved the performance a little bit.

In the experiments, we also included a simple BNN variant, where no filters are applied inside the permutohedral space, just splatting and slicing with the two layers BCL_a and BCL_b and adding the results. We will refer to this model as *BNN-Identity*, it corresponds to an image adaptive smoothing of the inputs $\mathbf{V}$. We found this filtering to have a positive effect and include it as a baseline in our experiments.

4.2. Spatial Network

The BNN was designed to propagate the information from the previous frames, respecting the scene and object motion. We then add a small spatial CNN with 3 layers, each with 32 filters of size 3×3, interleaved with ReLU non-linearities. The final result is then mapped to the desired output of V_t using a 1×1 convolution. The main role of this spatial CNN is to refine the information in frame t. Depending on the problem and the size of the available training data, other network designs are conceivable. We use the same network architecture shown in Fig. 3 for all the experiments to demonstrate the generality of VPNs.

5. Experiments

We evaluated VPN on three different propagation tasks: foreground masks, semantic labels and color information in videos. Our implementation runs in Caffe [35] using standard settings. We used Adam [38] stochastic optimization for training VPNs, multinomial-logistic loss for label propagation networks and Euclidean loss for training color propagation networks. Runtime computations were performed using a Nvidia TitanX GPU and a 6 core Intel i7-5820K CPU clocked at 3.30GHz machine. We will make available all the code and experimental results.

5.1. Video Object Segmentation

The task of class-agnostic video object segmentation aims to segment foreground objects in videos. Since the semantics of the foreground object is not pre-defined, this problem is usually addressed in a semi-supervised manner. The goal is to propagate a given foreground mask of the first frame to all the video frames. Object segmentation in videos is useful for several high level tasks such as video editing, summarization, rotoscoping etc.

Dataset We use the recently published DAVIS dataset [59] for experiments on this task. The DAVIS

	F-1	F-2	F-3	F-4	F-5	All
BNN-Identity	56.4	74.0	66.1	72.2	66.5	67.0
VPN-Stage1	58.2	77.7	70.4	76.0	68.1	70.1
VPN-Stage2	**60.9**	**78.7**	**71.4**	**76.8**	**69.0**	**71.3**

Table 1. **5-Fold Validation on DAVIS Video Segmentation Dataset.** Average IoU scores for different models on the 5 folds.

dataset consists of 50 high-quality (1080p resolution) unconstrained videos with number of frames in each video ranging from 25 to 104. All the frames come with high-quality per-pixel annotation of the foreground object. The videos for this dataset are carefully chosen to contain motion blur, occlusions, view-point changes and other occurrences of object segmentation challenges. For robust evaluation and to get results on all the dataset videos, we evaluate our technique using 5-fold cross-validation. We randomly divided the data into 5 folds, where in each fold, we used 35 images for training, 5 for validation and the remaining 10 for the testing. For the evaluation, we used the 3 metrics that are proposed in [59]: Intersection over Union (IoU) score, Contour accuracy ($\mathcal{F}$) score and temporal instability ($\mathcal{T}$) score. The widely used IoU score is defined as $TP/(TP+FN+FP)$, where TP: True positives; FN: False negatives and FP: False positives. Please refer to [59] for the definition of the contour accuracy and temporal instability scores. We are aware of some other datasets for this task such as JumpCut [26] and SegTrack [75], but we note that the number of videos in these datasets is too small for a learning based approach.

VPN and Results In this task, we only have access to foreground mask for the first frame V_1. For the ease of training VPN, we obtain initial set of predictions with *BNN-Identity*. We sequentially apply *BNN-Identity* at each frame and obtain an initial set of foreground masks for the entire video. These BNN-Identity propagated masks are then used as inputs to train a VPN to predict the refined masks at each frame. We refer to this VPN model as *VPN-Stage1*. Once VPN-Stage1 is trained, its refined training mask predictions are in-turn used as inputs to train another VPN model which we refer to as *VPN-Stage2*. This resulted in further refinement of foreground masks. Training further stages did not result in any improvements.

Following the recent work of [53] on video object segmentation, we used scaled features $\mathbf{f} = (x, y, Y, Cb, Cr, t)$ with YCbCr color features for bilateral filtering. To be comparable with the one of the fastest state-of-the-art technique [53], we do not use any optical flow information. First, we analyze the performance of BNN-Identity by changing the number of randomly sampled input points. Figure 4 shows how the segmentation IoU changes with the increase in the number of sampled points (out of 2 million points) from the previous frames. The IoU levels out after sampling 25% of points. For further computational efficiency, we used superpixel sampling instead of random sampling. Usage of superpixels reduced the IoU slightly (0.5%), while reducing the number of input points by a factor of 10 in comparison to a large number of randomly sampled points. We used 12000 SLIC [1] superpixels from each frame computed using the fast GPU implementation from [63]. For predictions at each frame, we input mask probabilities of previous 9 frames into VPN as we observe

	IoU↑	$\mathcal{F}\uparrow$	$\mathcal{T}\downarrow$	*Runtime*(s)
BNN-Identity	67.0	67.1	36.3	0.21
VPN-Stage1	70.1	68.4	30.1	0.48
VPN-Stage2	71.3	68.9	30.2	0.75
With pre-trained models				
DeepLab	57.0	49.9	47.8	0.15
VPN-DeepLab	**75.0**	**72.4**	29.5	0.63
OFL [76]	71.1	67.9	22.1	>60
BVS [53]	66.5	65.6	31.6	0.37
NLC [25]	64.1	59.3	35.6	20
FCP [60]	63.1	54.6	28.5	12
JMP [26]	60.7	58.6	**13.2**	12
HVS [29]	59.6	57.6	29.7	5
SEA [62]	55.6	53.3	13.7	6

Table 2. **Results of Video Object Segmentation on DAVIS dataset.** Average IoU score, contour accuracy ($\mathcal{F}$), temporal instability ($\mathcal{T}$) scores, and average runtimes (in seconds) per frame for different VPN models along with recent published techniques for this task. VPN runtimes also include superpixel computation (10ms). Runtimes of other methods are taken from [53, 60, 76] and only indicative and are not directly comparable to our runtimes. Runtime of VPN-Stage1 includes the runtime of BNN-Identity which is in-turn included in the runtime of VPN-Stage2. Runtime of VPN-DeepLab model includes the runtime of DeepLab.

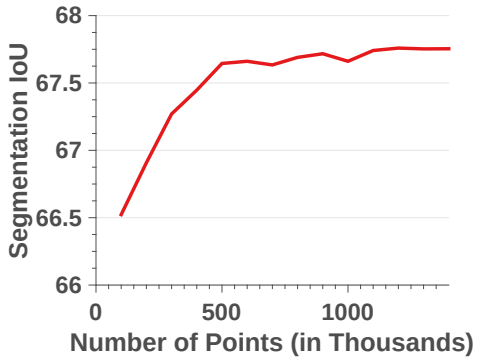

Figure 4. **Random Sampling of Input Points vs. IoU.** The effect of randomly sampling points from input video frames on object segmentation IoU of BNN-Identity on DAVIS dataset. The points sampled are out of ≈2 million points from the previous 5 frames.

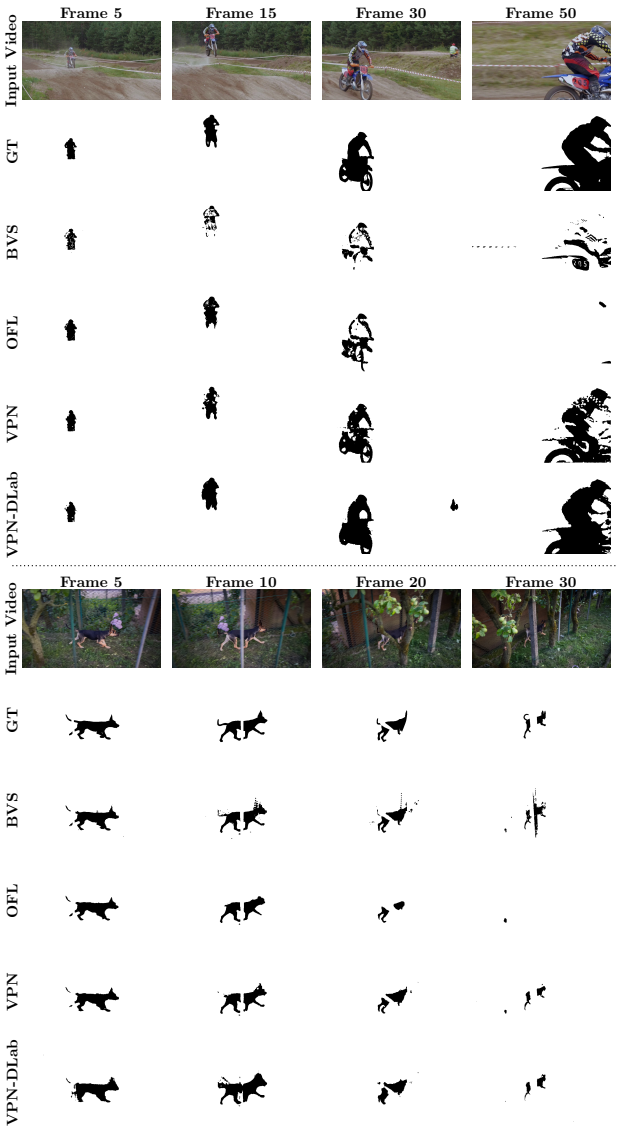

Figure 5. **Video Object Segmentation.** Shown are the different frames in example videos with the corresponding ground truth (GT) masks, predictions from BVS [53], OFL [76], VPN (VPN-Stage2) and VPN-DLab (VPN-DeepLab) models.

no significant improvements with more frames. We set α to 0.5 and the feature scales for bilateral filtering are presented in Tab. A.1.

Table 1 shows the IoU scores for each of the 5 folds and Tab. 2 shows the overall scores and runtimes of different VPN models along with the best performing segmentation techniques. The performance improved consistently across all 5 folds with the addition of new VPN stages. BNN-Identity already performed reasonably well. And with 1-stage and 2-stage VPNs, we outperformed the present fastest BVS method [53] by a significant margin on all the performance measures of IoU, contour accuracy and temporal instability scores, while being comparable in runtime. We perform marginally better than OFL method [76] while being at least 80$\times$ faster and OFL relies on optical flow whereas we obtain similar performance without using any optical flow. Further, VPN has the advantage of doing online processing as it looks only at previous frames whereas BVS processes entire video at once. One can obtain better VPN performance with using better superpixels and also incorporating optical flow, but this increases runtime as well. Figure 5 shows some qualitative results and more are present in Figs. A.1, A.2. A couple of failure cases are shown in Fig. A.3. Visual results indicate that learned VPN is able to retain foreground masks even with large variations in viewpoint and object size.

Augmenation of Pre-trained Models: One of the main advantages of the proposed VPN architecture is that it is end-to-end trainable and can be easily integrated into other deep neural network architectures. To demonstrate this, we augmented VPN architecture with standard DeepLab segmentation architecture from [17]. We replaced the last classification layer of DeepLab-LargeFOV model from [17] to output 2 classes (foreground and background) in our case and bi-linearly up-sampled the resulting low-resolution probability map to the original image dimension. 5-fold fine-tuning of the DeepLab model on DAVIS dataset resulted in the IoU of 57.0 and other scores are shown in Tab. 2. Then, we combine the VPN and DeepLab models in the following way: The outputs from the DeepLab network and the bilateral network are concatenated and then passed on to the spatial network. In other words, the bilateral network propagates label information from previous frames to the present frame, whereas the DeepLab network does the prediction for the present frame. The results of both are then combined and refined by the spatial network in the VPN architecture. We call this 'VPN-DeepLab' model. We trained this model end-to-end and observed big improvements in performance. As shown in Tab. 2, the VPN-DeepLab model has the IoU score of 75.0 and contour accuracy score of 72.4 resulting in significant improvements over the published results. Since DeepLab has also fast runtime, the total runtime of VPN-DeepLab is only 0.63s which makes this also one of the fastest video segmentation systems. A couple of visual results of VPN-DeepLab model are shown in Fig. 5 and more are present in Figs. A.1, A.2 and A.3.

5.2. Semantic Video Segmentation

A semantic video segmentation assigns a semantic label to every video pixel. Since the semantics between adjacent frames does not change radically, intuitively, propagating semantic information across frames should improve the segmentation quality of each frame. Unlike mask propagation in the previous section where the ground-truth mask for the first frame is given, we approach semantic video segmentation in a fully automatic fashion. Specifically, we start with

	IoU	*Runtime*(s)
CNN from [80]	65.3	0.38
+ FSO-CRF [43]	66.1	>10
+ BNN-Identity	65.3	0.31
+ BNN-Identity-Flow	65.5	0.33
+ VPN (Ours)	66.5	0.35
+ VPN-Flow (Ours)	**66.7**	0.37
CNN from [65]	68.9	0.30
+ VPN-Flow (Ours)	**69.5**	0.38

Table 3. **Results of Semantic Segmentation on the CamVid Dataset.** Average IoU and runtimes (in seconds) per frame of different models on *test* split. Runtimes exclude CNN computations which are shown separately. VPN and BNN-Identity runtimes include superpixel computation which takes up large portion of computation time (0.23s).

the unary predictions of standard CNNs and use VPN for propagating semantics across the frames.

Dataset We use the CamVid dataset [11] that contains 4 high quality videos captured at 30Hz while the semantically labeled 11-class ground truth is provided at 1Hz. While the original dataset comes at a resolution of 960×720, similar to previous works [80, 43], we operate on a resolution of 640×480. We use the same splits proposed in [70] resulting in 367, 100 and 233 frames with ground truth for training, validation and testing. Following common practice, we report the IoU scores for evaluation.

VPN and Results Since we already have CNN predictions for every frame, we train a VPN that takes the CNN predictions of previous *and* present frames as input and predicts the refined predictions for the present frame. We compare with the state-of-the-art CRF approach for this problem [43] which we refer to as FSO-CRF. Following [43], we also experimented with optical flow in our framework and refer that model as *VPN-Flow*. We used the fast optical flow method that uses dense inverse search [42] to compute flows and modify the positional features of previous frames. We used the superpixels method of Dollar et al. [20] for this dataset as gSLICr [63] has introduced artifacts.

We experimented with predictions from two different CNNs: One is with dilated convolutions [80] (CNN-1) and another one [65] (CNN-2) is trained with the additional data obtained from a video game, which is the present state-of-the-art on this dataset. For CNN-1 and CNN-2, using 2 and 3 previous frames respectively as input to VPN is found to be optimal. Other parameters of the bilateral network are presented in Tab. A.1. Table 3 shows quantitative results on this dataset. Using BNN-Identity only slightly improved the CNN performance whereas training the entire VPN significantly improved the CNN performance by over 1.2% IoU, with both VPN and VPN-Flow networks. Moreover, VPN is at least 25× faster, and simpler to use compared to the optimization based FSO-CRF which relies on LDOF optical flow [13], long-term tacks [71] and edges [21]. We further improved the performance of the state-of-the-art CNN [65] with the use of VPN-Flow model. Using better optical flow estimation might give even better results. Figure 6 shows some qualitative results and more are presented in Fig. A.4.

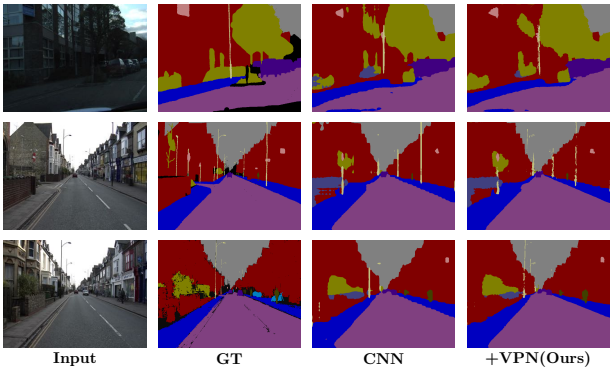

Figure 6. **Semantic Video Segmentation.** Input video frames and the corresponding ground truth (GT) segmentation together with the predictions of CNN [80] and with VPN-Flow.

5.3. Video Color Propagation

We also evaluate VPNs on a different kind of information and experimented with propagating color information in a grayscale video. Given the color image for the first video frame, the task is to propagate the color to the entire video. Note that this task is fundamentally different from automatic colorization of images for which recent CNN based based methods have become popular. For experiments on this task, we again used the DAVIS dataset [59] with the first 25 frames from each video. We randomly divided the dataset into 30 train, 5 validation and 15 test videos.

We work with YCbCr representation of images and propagate CbCr values from previous frames with pixel intensity, position and time features as guidance for VPN. The same strategy as in object segmentation is used, where an initial set of color propagated results was obtained with BNN-Identity and then used to trained a VPN-Stage1 model. Training further VPN stages did not improve the performance. Table 4 shows the PSNR results. We use 300K randomly sampled points from previous 3 frames as input to the VPN network. We also show a baseline result of [46] that does graph based optimization and uses optical flow. We used fast DIS optical flow [42] in the baseline method [46] and we did not observe significant differences with using LDOF optical flow [13]. Figure 7 shows a visual result with more in Fig. A.5. From the results, VPN works reliably better than [46] while being 20× faster. The method of [46] relies heavily on optical flow and so the color drifts away with incorrect flow. We observe that our method also

	PSNR	*Runtime*(s)
BNN-Identity	27.89	0.29
VPN-Stage1	**28.15**	0.90
Levin et al. [46]	27.11	19

Table 4. **Results of Video Color Propagation.** Average PSNR results and runtimes of different methods for video color propagation on images from DAVIS dataset.

bleeds color in some regions especially when there are large viewpoint changes. We could not compare against recent video color propagation techniques such as [31, 68] as their codes are not available online. This application shows general applicability of VPNs in propagating different kinds of information.

Figure 7. **Video Color Propagation.** Input grayscale video frames and corresponding ground-truth (GT) color images together with color predictions of Levin et al. [46] and VPN-Stage1 models.

6. Conclusion

We proposed a fast, scalable and generic neural network based learning approach for propagating information across video frames. The video propagation network uses bilateral network for long-range video-adaptive propagation of information from previous frames to the present frame which is then refined by a standard spatial network. Experiments on diverse tasks show that VPNs, despite being generic, outperformed the current state-of-the-art task-specific methods. At the core of our technique is the exploitation and modification of learnable bilateral filtering for the use in video processing. We used a simple and fixed network architecture for all the tasks for showcasing the generality of the approach. Depending on the type of problems and the availability of data, using more filters and deeper layers would result in better performance. In this work, we manually tuned the feature scales which could be amendable to learning. Finding optimal yet fast-to-compute bilateral features for videos together with the learning of their scales is an important future research direction.

Acknowledgments We thank Vibhav Vineet for providing the trained image segmentation CNN models for CamVid dataset.

References

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and S. Susstrunk. SLIC superpixels compared to state-of-the-art superpixel methods. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(11):2274–2282, 2012. 6

[2] A. Adams, J. Baek, and M. A. Davis. Fast high-dimensional filtering using the permutohedral lattice. In *Computer Graphics Forum*, volume 29, pages 753–762. Wiley Online Library, 2010. 3, 4

[3] A. Adams, N. Gelfand, J. Dolson, and M. Levoy. Gaussian kd-trees for fast high-dimensional filtering. *ACM Transactions on Graphics (ToG)*, 28(3):21, 2009. 3

[4] A. Agarwala, A. Hertzmann, D. H. Salesin, and S. M. Seitz. Keyframe-based tracking for rotoscoping and animation. *ACM Transactions on Graphics (ToG)*, 23(3):584–591, 2004. 2

[5] V. Aurich and J. Weule. Non-linear Gaussian filters performing edge preserving diffusion. In *DAGM*, pages 538–545. Springer, 1995. 2, 3

[6] X. Bai, J. Wang, D. Simons, and G. Sapiro. Video snapcut: robust video object cutout using localized classifiers. *ACM Transactions on Graphics (TOG)*, 28(3):70, 2009. 2

[7] J. T. Barron, A. Adams, Y. Shih, and C. Hernández. Fast bilateral-space stereo for synthetic defocus. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4466–4474, 2015. 2

[8] J. T. Barron and B. Poole. The fast bilateral solver. In *European Conference on Computer Vision*. Springer, 2016. 2

[9] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(11):1222–1239, 2001. 2

[10] Y. Y. Boykov and M.-P. Jolly. Interactive graph cuts for optimal boundary & region segmentation of objects in nd images. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 1, pages 105–112. IEEE, 2001. 2

[11] G. J. Brostow, J. Fauqueur, and R. Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2):88–97, 2009. 8

[12] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla. Segmentation and recognition using structure from motion point clouds. In *European Conference on Computer Vision*, pages 44–57. Springer, 2008. 2

[13] T. Brox, C. Bregler, and J. Malik. Large displacement optical flow. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 41–48. IEEE, 2009. 8

[14] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-shot video object segmentation. *arXiv preprint arXiv:1611.05198*, 2016. 2

[15] J.-H. R. Chang and Y.-C. F. Wang. Propagated image filtering. In *Computer Vision and Pattern Recognition (CVPR), IEEE Conference on*, pages 10–18. IEEE, 2015. 2

[16] A. Y. Chen and J. J. Corso. Temporally consistent multi-class video-object segmentation with the video graph-shifts algorithm. In *IEEE Winter Conference on Applications of Computer Vision*, pages 614–621. IEEE, 2011. 2

[17] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014. 3, 7

[18] Y.-Y. Chuang, A. Agarwala, B. Curless, D. H. Salesin, and R. Szeliski. Video matting of complex scenes. *ACM Transactions on Graphics (ToG)*, 21(3):243–248, 2002. 2

[19] R. de Nijs, S. Ramos, G. Roig, X. Boix, L. Van Gool, and K. Kühnlenz. On-line semantic perception using uncertainty. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 4185–4191. IEEE, 2012. 2, 3

[20] P. Dollár and C. L. Zitnick. Structured forests for fast edge detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 2013. 8

[21] P. Dollár and C. L. Zitnick. Fast edge detection using structured forests. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(8):1558–1570, 2015. 8

[22] J. Domke. Learning graphical model parameters with approximate marginal inference. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(10):2454–2467, 2013. 2

[23] R. Dondera, V. Morariu, Y. Wang, and L. Davis. Interactive video segmentation using occlusion boundaries and temporally coherent superpixels. In *IEEE Winter Conference on Applications of Computer Vision*, pages 784–791. IEEE, 2014. 2

[24] A. Ess, T. Mueller, H. Grabner, and L. J. Van Gool. Segmentation-based urban traffic scene understanding. In *British Machine Vision Conference*, 2009. 2

[25] A. Faktor and M. Irani. Video segmentation by non-local consensus voting. In *BMVC*, volume 2, page 6, 2014. 2, 6

[26] Q. Fan, F. Zhong, D. Lischinski, D. Cohen-Or, and B. Chen. Jumpcut: non-successive mask transfer and interpolation for video cutout. *ACM Transactions on Graphics (ToG)*, 34(6):195, 2015. 2, 6

[27] R. Gadde, V. Jampani, M. Kiefel, D. Kappler, and P. Gehler. Superpixel convolutional networks using bilateral inceptions. In *European Conference on Computer Vision*. Springer, 2016. 3

[28] E. S. Gastal and M. M. Oliveira. Domain transform for edge-aware image and video processing. *ACM Transactions on Graphics (ToG)*, 30(4):69, 2011. 3

[29] M. Grundmann, V. Kwatra, M. Han, and I. Essa. Efficient hierarchical graph-based video segmentation. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 2141–2148. IEEE, 2010. 6

[30] K. He, J. Sun, and X. Tang. Guided image filtering. In *European Conference on Computer Vision*, pages 1–14. Springer, 2010. 2

[31] J.-H. Heu, D.-Y. Hyun, C.-S. Kim, and S.-U. Lee. Image and video colorization based on prioritized source propagation. In *2009 16th IEEE International Conference on Image Processing (ICIP)*, pages 465–468. IEEE, 2009. 8

[32] T.-W. Hui, C. C. Loy, and X. Tang. Depth map super-resolution by deep multi-scale guidance. In *European Conference on Computer Vision*, pages 353–369. Springer, 2016. 2

[33] S. D. Jain and K. Grauman. Supervoxel-consistent foreground propagation in video. In *European Conference on Computer Vision*, pages 656–671. Springer, 2014. 2

[34] V. Jampani, M. Kiefel, and P. V. Gehler. Learning sparse high dimensional filters: Image filtering, dense crfs and bilateral neural networks. In *Computer Vision and Pattern Recognition, IEEE Conference on*, June 2016. 2, 3, 4, 5

[35] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014. 5

[36] A. Khoreva, F. Perazzi, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. *arXiv preprint arXiv:1612.02646*, 2016. 2

[37] M. Kiefel, V. Jampani, and P. V. Gehler. Permutohedral lattice CNNs. *International Conference on Learning Representations Workshop*, 2015. 2, 3, 5

[38] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 2015. 5

[39] P. Kohli and P. H. Torr. Dynamic graph cuts for efficient inference in markov random fields. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 29(12):2079–2088, 2007. 2

[40] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele. Joint bilateral upsampling. *ACM Transactions on Graphics (ToG)*, 26(3):96, 2007. 2, 4

[41] P. Krähenbühl and V. Koltun. Efficient inference in fully connected CRFs with Gaussian edge potentials. In *Advances in neural information processing systems*, 2011. 2

[42] T. Kroeger, R. Timofte, D. Dai, and L. Van Gool. Fast optical flow using dense inverse search. In *European Conference on Computer Vision*. Springer, 2016. 8

[43] A. Kundu, V. Vineet, and V. Koltun. Feature space optimization for semantic video segmentation. In *Computer Vision and Pattern Recognition, IEEE Conference on*, 2016. 2, 3, 8

[44] M. Lang, O. Wang, T. O. Aydin, A. Smolic, and M. H. Gross. Practical temporal consistency for image-based graphics applications. *ACM Transactions on Graphics (ToG)*, 31(4):34–1, 2012. 2

[45] Y. J. Lee, J. Kim, and K. Grauman. Key-segments for video object segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1995–2002. IEEE, 2011. 2

[46] A. Levin, D. Lischinski, and Y. Weiss. Colorization using optimization. *ACM Transactions on Graphics (ToG)*, 23(3):689–694, 2004. 2, 8, 9, 17

[47] A. Levin, D. Lischinski, and Y. Weiss. A closed-form solution to natural image matting. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 30(2):228–242, 2008. 2

[48] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg. Video segmentation by tracking many figure-ground segments. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2192–2199, 2013. 2

[49] Y. Li, J.-B. Huang, N. Ahuja, and M.-H. Yang. Deep joint image filtering. In *European Conference on Computer Vision*, pages 154–169. Springer, 2016. 2

[50] Y. Li, J. Sun, and H.-Y. Shum. Video object cut and paste. *ACM Transactions on Graphics (ToG)*, 24(3):595–600, 2005. 2

[51] Y. Li and R. Zemel. Mean-field networks. *arXiv preprint arXiv:1410.5884*, 2014. 2

[52] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 3431–3440, 2015. 3

[53] N. Märki, F. Perazzi, O. Wang, and A. Sorkine-Hornung. Bilateral space video segmentation. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 743–751, 2016. 2, 6, 7, 13, 14, 15

[54] O. Miksik, D. Munoz, J. A. Bagnell, and M. Hebert. Efficient temporal consistency for streaming video scene analysis. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 133–139. IEEE, 2013. 2, 3

[55] A. Papazoglou and V. Ferrari. Fast object segmentation in unconstrained video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1777–1784, 2013. 2

[56] S. Paris. Edge-preserving smoothing and mean-shift segmentation of video streams. In *European Conference on Computer Vision*, pages 460–473. Springer, 2008. 4

[57] S. Paris and F. Durand. A fast approximation of the bilateral filter using a signal processing approach. In *European Conference on Computer Vision*, pages 568–580. Springer, 2006. 3

[58] S. Paris, P. Kornprobst, J. Tumblin, and F. Durand. *Bilateral filtering: Theory and applications*. Now Publishers Inc, 2009. 3

[59] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. V. Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Computer Vision and Pattern Recognition, IEEE Conference on*, 2016. 5, 6, 8

[60] F. Perazzi, O. Wang, M. Gross, and A. Sorkine-Hornung. Fully connected object proposals for video segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3227–3234, 2015. 2, 6

[61] B. L. Price, B. S. Morse, and S. Cohen. Livecut: Learning-based interactive video segmentation by evaluation of multiple propagated cues. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 779–786. IEEE, 2009. 2

[62] S. A. Ramakanth and R. V. Babu. Seamseg: Video object segmentation using patch seams. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 376–383. IEEE, 2014. 6

[63] C. Y. Ren, V. A. Prisacariu, and I. D. Reid. gslicr: Slic superpixels at over 250hz. *ArXiv e-prints*, (1509.04232), 2015. 6, 8

[64] M. Reso, B. Scheuermann, J. Jachalsky, B. Rosenhahn, and J. Ostermann. Interactive segmentation of high-resolution video content using temporally coherent superpixels and graph cut. In *International Symposium on Visual Computing*, pages 281–292. Springer, 2014. 2

[65] S. R. Richter, V. Vineet, S. Roth, and V. Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision*, pages 102–118. Springer, 2016. 8, 13

[66] A. G. Schwing and R. Urtasun. Fully connected deep structured networks. *arXiv preprint arXiv:1503.02351*, 2015. 2

[67] E. Shelhamer, K. Rakelly, J. Hoffman, and T. Darrell. Clockwork convnets for video semantic segmentation. *arXiv preprint arXiv:1608.03609*, 2016. 3

[68] B. Sheng, H. Sun, M. Magnor, and P. Li. Video colorization using parallel optimization in feature space. *IEEE Transactions on Circuits and Systems for Video Technology*, 24(3):407–417, 2014. 8

[69] J. Shi and J. Malik. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8):888–905, 2000. 2

[70] P. Sturgess, K. Alahari, L. Ladicky, and P. H. Torr. Combining appearance and structure from motion features for road scene understanding. In *British Machine Vision Conference*. BMVA, 2009. 2, 8

[71] N. Sundaram, T. Brox, and K. Keutzer. Dense point trajectories by gpu-accelerated large displacement optical flow. In *European Conference on Computer Vision*, pages 438–451. Springer, 2010. 8

[72] B. Taylor, V. Karasev, and S. Soattoc. Causal video object segmentation from persistence of occlusions. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 4268–4276. IEEE, 2015. 2

[73] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 839–846. IEEE, 1998. 2, 3

[74] S. Tripathi, S. Belongie, Y. Hwang, and T. Nguyen. Semantic video segmentation: Exploring inference efficiency. In *2015 International SoC Design Conference (ISOCC)*, pages 157–158. IEEE, 2015. 2

[75] D. Tsai, M. Flagg, A. Nakazawa, and J. M. Rehg. Motion coherent tracking using multi-label mrf optimization. *International journal of computer vision*, 100(2):190–202, 2012. 6

[76] Y.-H. Tsai, M.-H. Yang, and M. J. Black. Video segmentation via object flow. In *Computer Vision and Pattern Recognition, IEEE Conference on*, 2016. 2, 6, 7, 13, 14, 15

[77] J. Wang, P. Bhat, R. A. Colburn, M. Agrawala, and M. F. Cohen. Interactive video cutout. *ACM Transactions on Graphics (ToG)*, 24(3):585–594, 2005. 2

[78] W. Wang, J. Shen, and F. Porikli. Saliency-aware geodesic video object segmentation. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 3395–3402, 2015. 2

[79] L. Xu, J. S. Ren, Q. Yan, R. Liao, and J. Jia. Deep edge-aware filters. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 1669–1678, 2015. 2

[80] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. *International Conference on Learning Representations*, 2016. 8, 13, 16

[81] D. Zhang, O. Javed, and M. Shah. Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *Computer Vision and Pattern Recognition, IEEE Conference on*, pages 628–635, 2013. 2

[82] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. Torr. Conditional random fields as recurrent neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015. 2, 3

A. Parameters and Additional Results

In this appendix, we present experiment protocols and additional qualitative results for experiments on video object segmentation, semantic video segmentation and video color propagation. Table A.1 shows the feature scales and other parameters used in different experiments. Figures A.1, A.2 show some qualitative results on video object segmentation with some failure cases in Fig. A.3. Figure A.4 shows some qualitative results on semantic video segmentation and Fig. A.5 shows results on video color propagation.

Experiment	Feature Type	Feature Scale-1, Λ_a	Feature Scale-2, Λ_b	α	Input Frames	Loss Type
Video Object Segmentation	(x, y, Y, Cb, Cr, t)	(0.02,0.02,0.07,0.4,0.4,0.01)	(0.03,0.03,0.09,0.5,0.5,0.2)	0.5	9	Logistic
Semantic Video Segmentation with CNN1 [80]-NoFlow	(x, y, R, G, B, t)	(0.08,0.08,0.2,0.2,0.2,0.04)	(0.11,0.11,0.2,0.2,0.2,0.04)	0.5	3	Logistic
with CNN1 [80]-Flow	$(x+u_x, y+u_y, R, G, B, t)$	(0.11,0.11,0.14,0.14,0.14,0.03)	(0.08,0.08,0.12,0.12,0.12,0.01)	0.65	3	Logistic
with CNN2 [65]-Flow	$(x+u_x, y+u_y, R, G, B, t)$	(0.08,0.08,0.2,0.2,0.2,0.04)	(0.09,0.09,0.25,0.25,0.25,0.03)	0.5	4	Logistic
Video Color Propagation	(x, y, I, t)	(0.04,0.04,0.2,0.04)	No second kernel	1	4	MSE

Table A.1. **Experiment Protocols.** Experiment protocols for the different experiments presented in this work. **Feature Types**: Feature spaces used for the bilateral convolutions, with position (x, y) and color $(R, G, B$ or $Y, Cb, Cr)$ features $\in [0, 255]$. u_x, u_y denotes optical flow with respect to the present frame and I denotes grayscale intensity. **Feature Scales** (Λ_a, Λ_b): Cross-validated scales for the features used. α: Exponential time decay for the input frames. **Input Frames**: Number of input frames for VPN. **Loss Type**: Type of loss used for back-propagation. "MSE" corresponds to Euclidean mean squared error loss and "Logistic" corresponds to multinomial logistic loss.

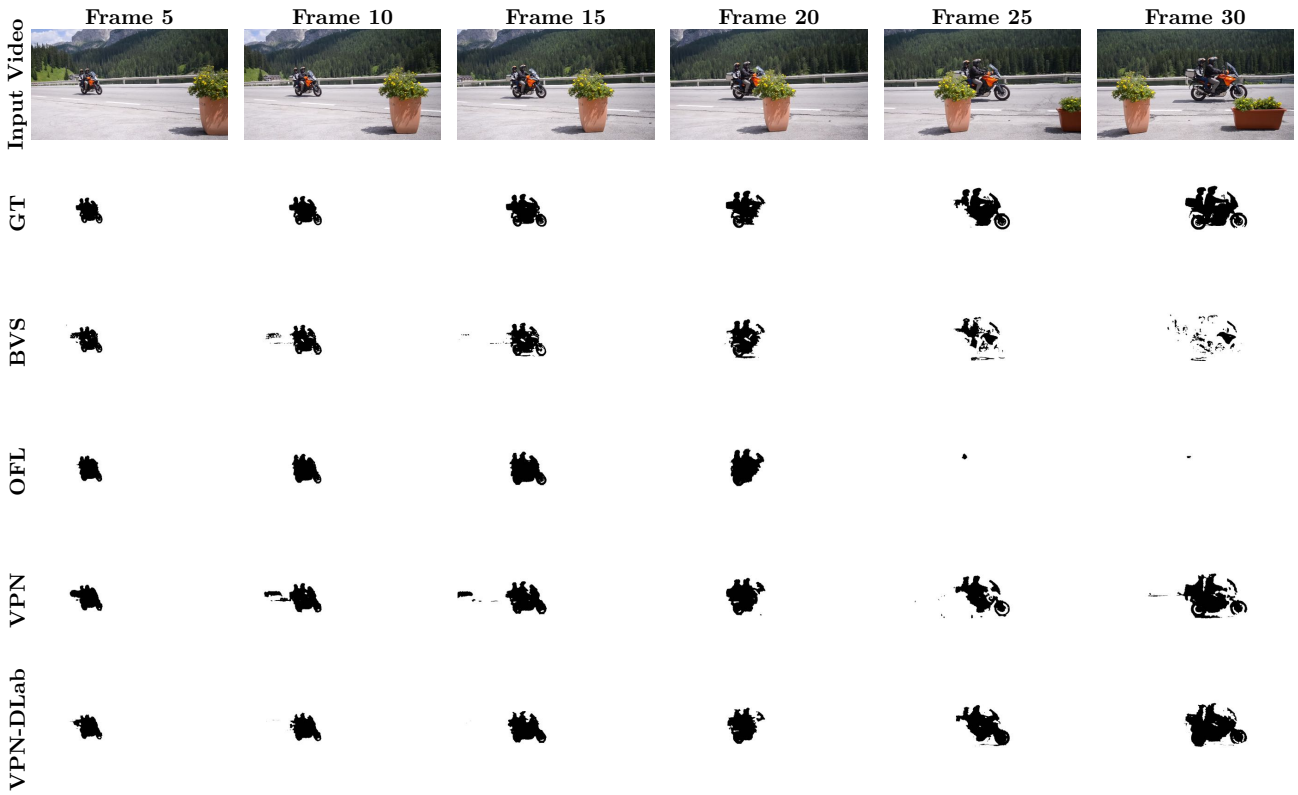

Figure A.1. **Video Object Segmentation.** Shown are the different frames in example videos with the corresponding ground truth (GT) masks, predictions from BVS [53], OFL [76], VPN (VPN-Stage2) and VPN-DLab (VPN-DeepLab) models.

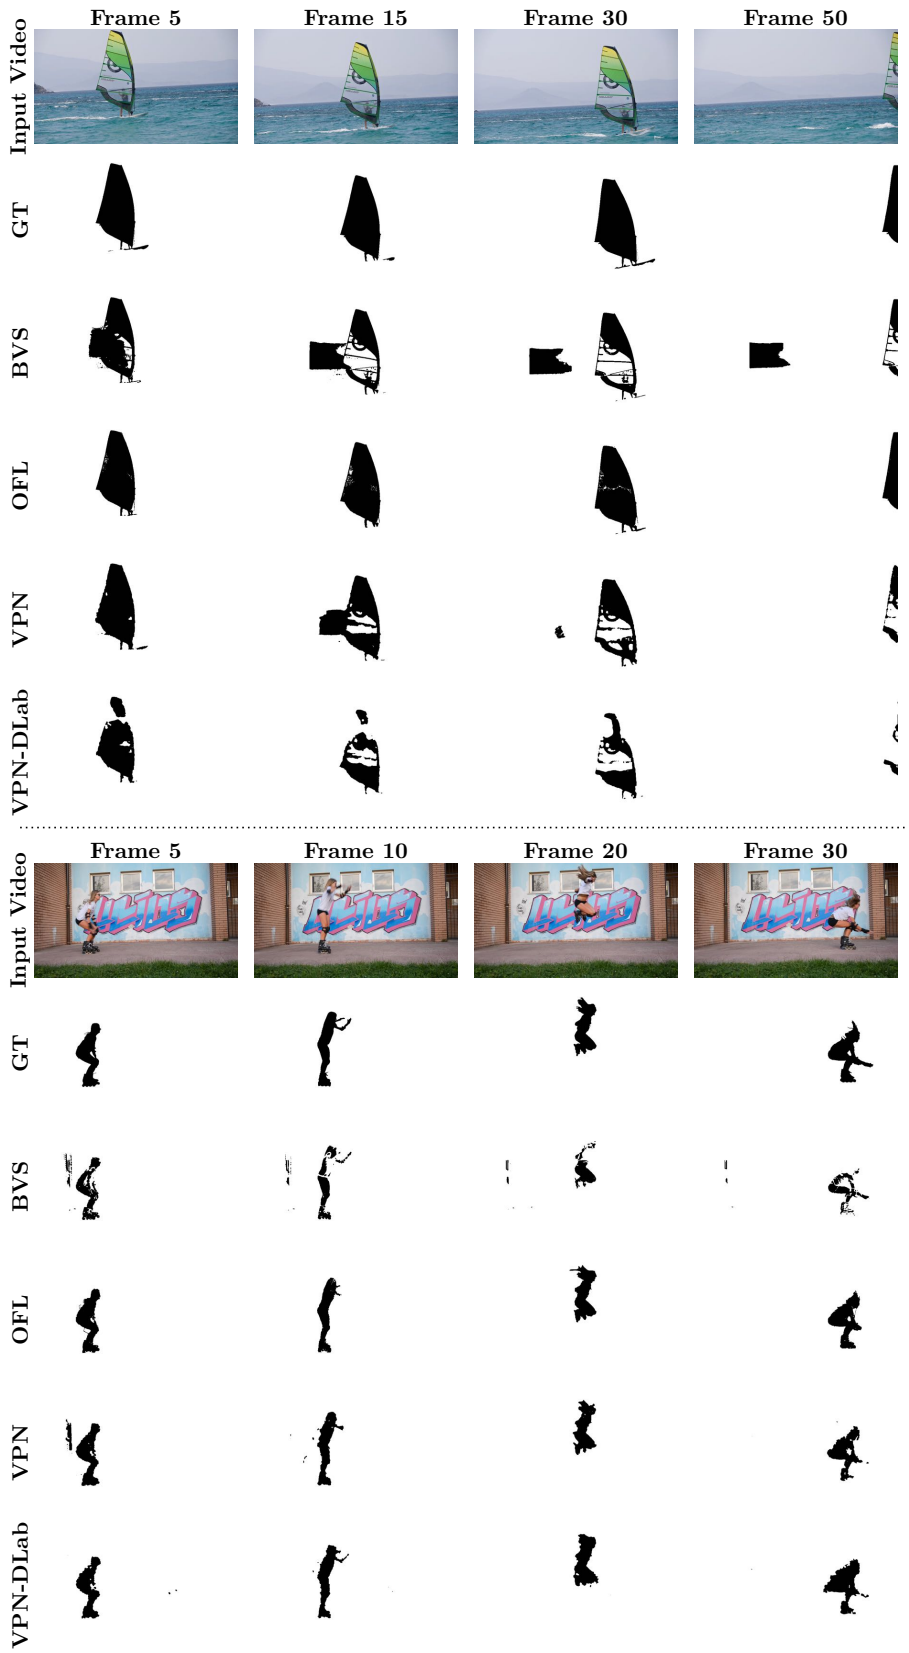

Figure A.2. **Video Object Segmentation.** Shown are the different frames in example videos with the corresponding ground truth (GT) masks, predictions from BVS [53], OFL [76], VPN (VPN-Stage2) and VPN-DLab (VPN-DeepLab) models.

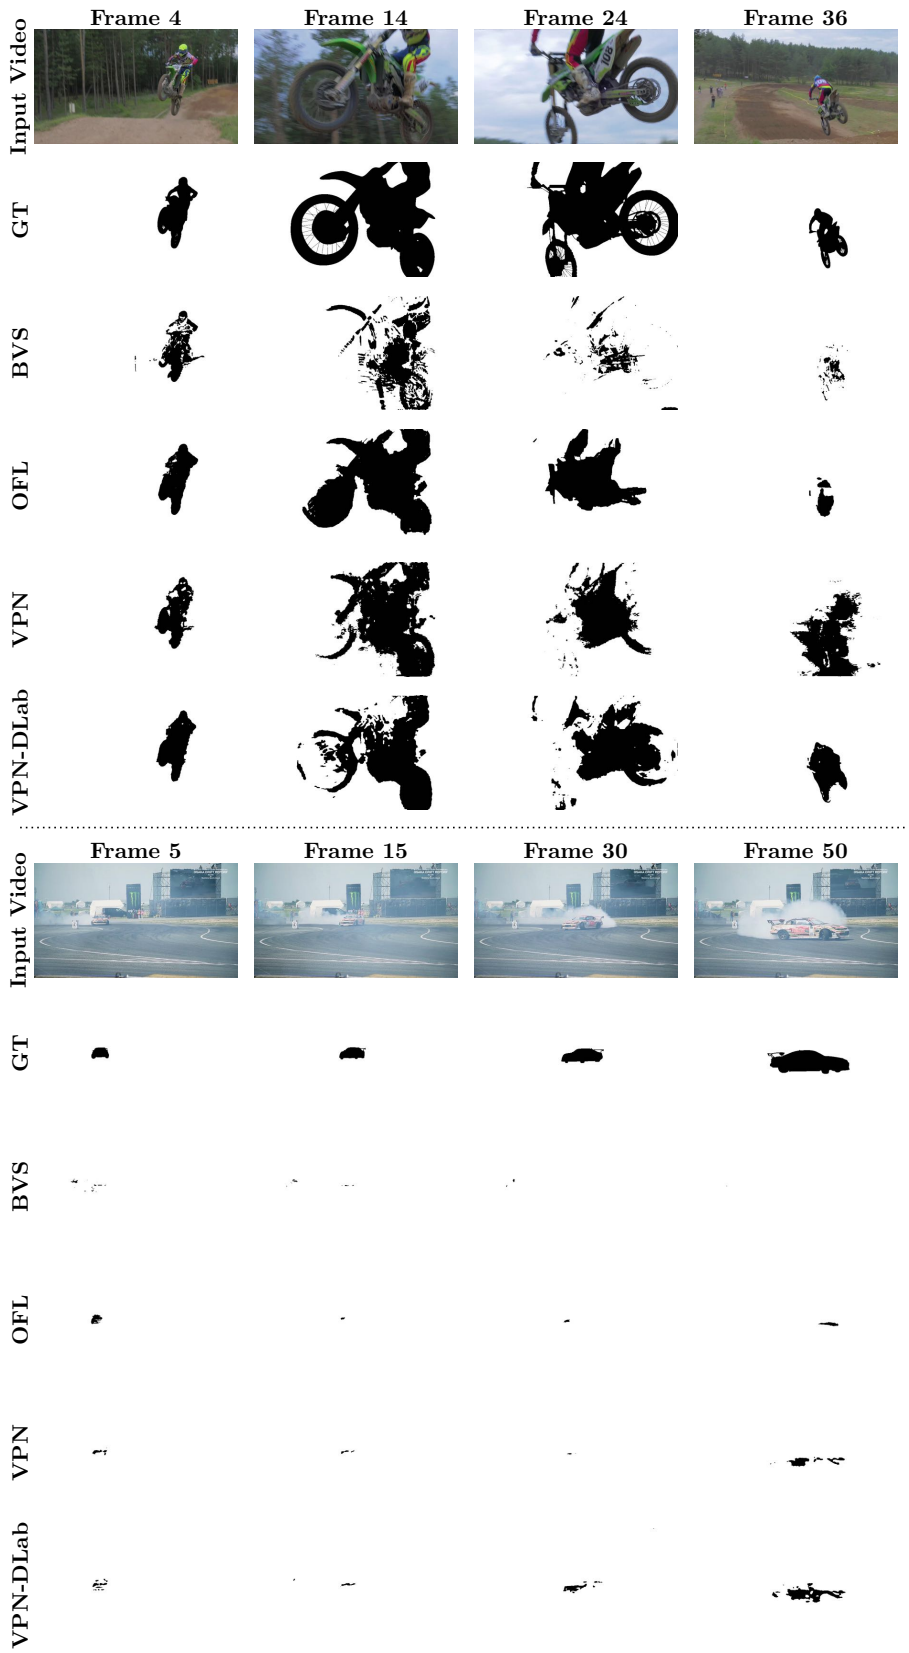

Figure A.3. **Failure Cases for Video Object Segmentation.** Shown are the different frames in example videos with the corresponding ground truth (GT) masks, predictions from BVS [53], OFL [76], VPN (VPN-Stage2) and VPN-DLab (VPN-DeepLab) models.

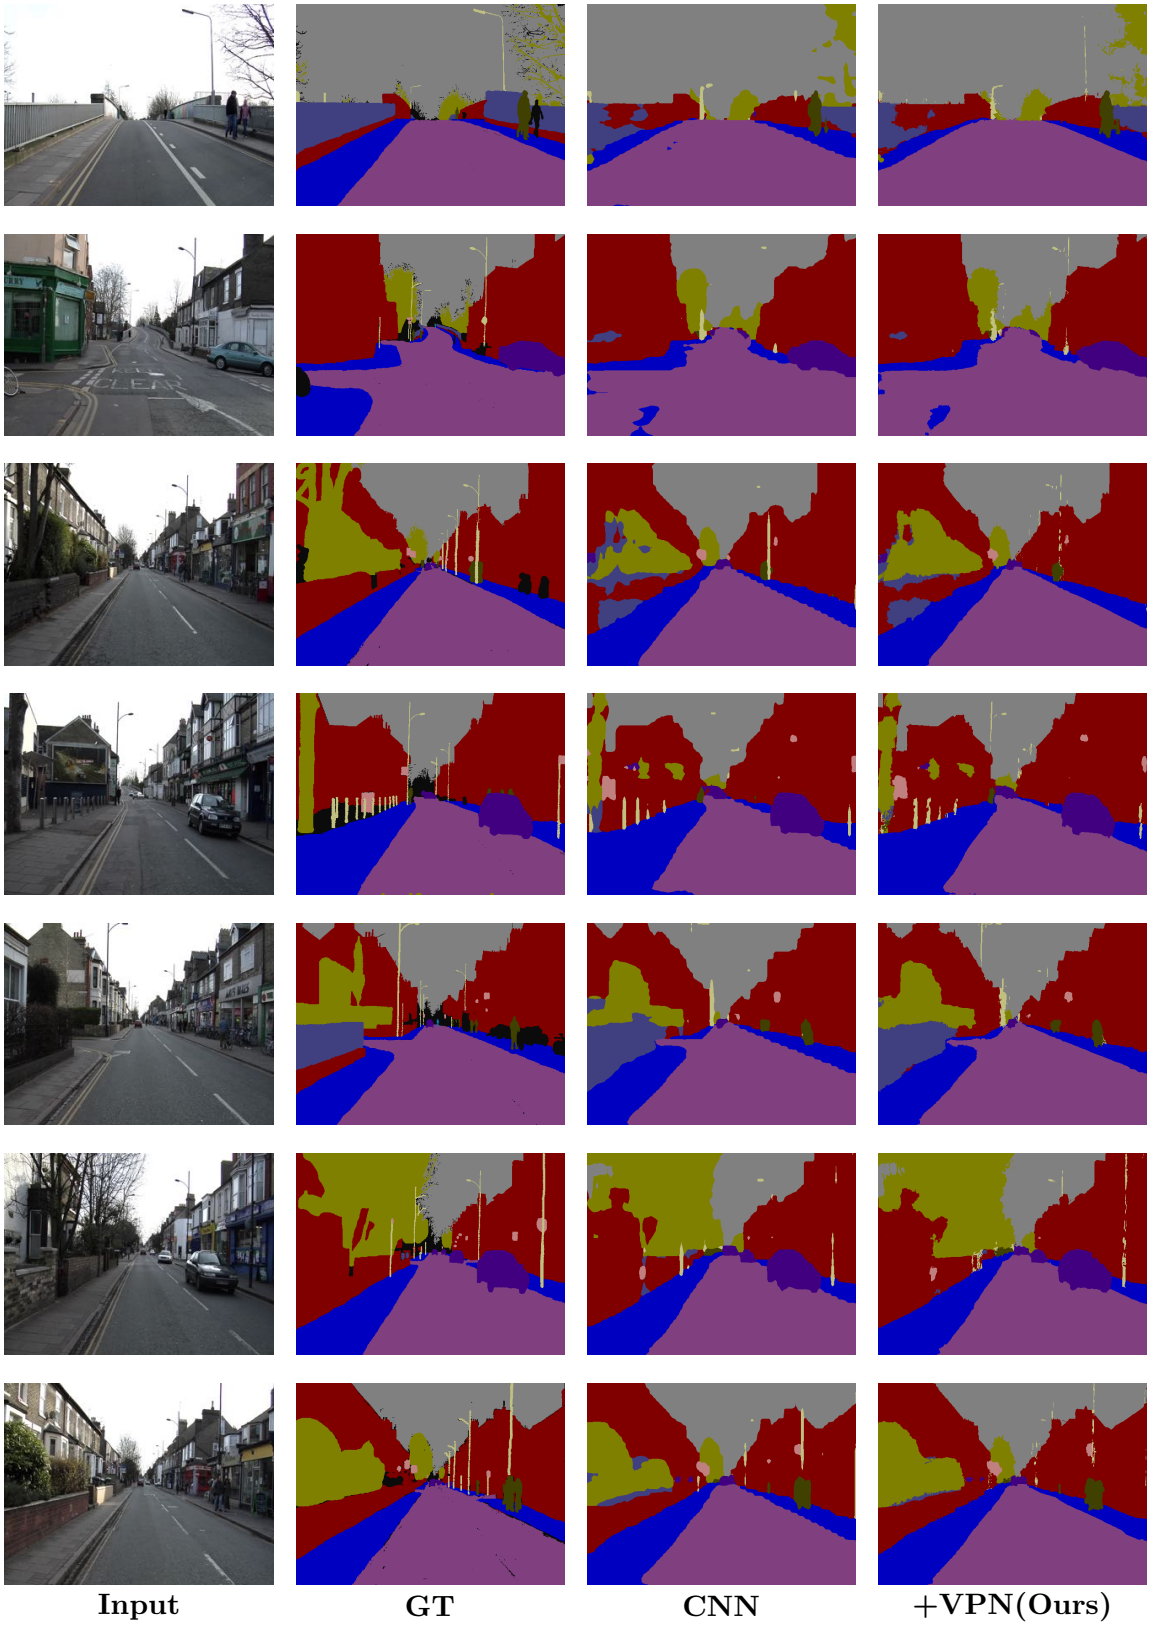

Figure A.4. **Semantic Video Segmentation.** Input video frames and the corresponding ground truth (GT) segmentation together with the predictions of CNN [80] and with VPN-Flow.

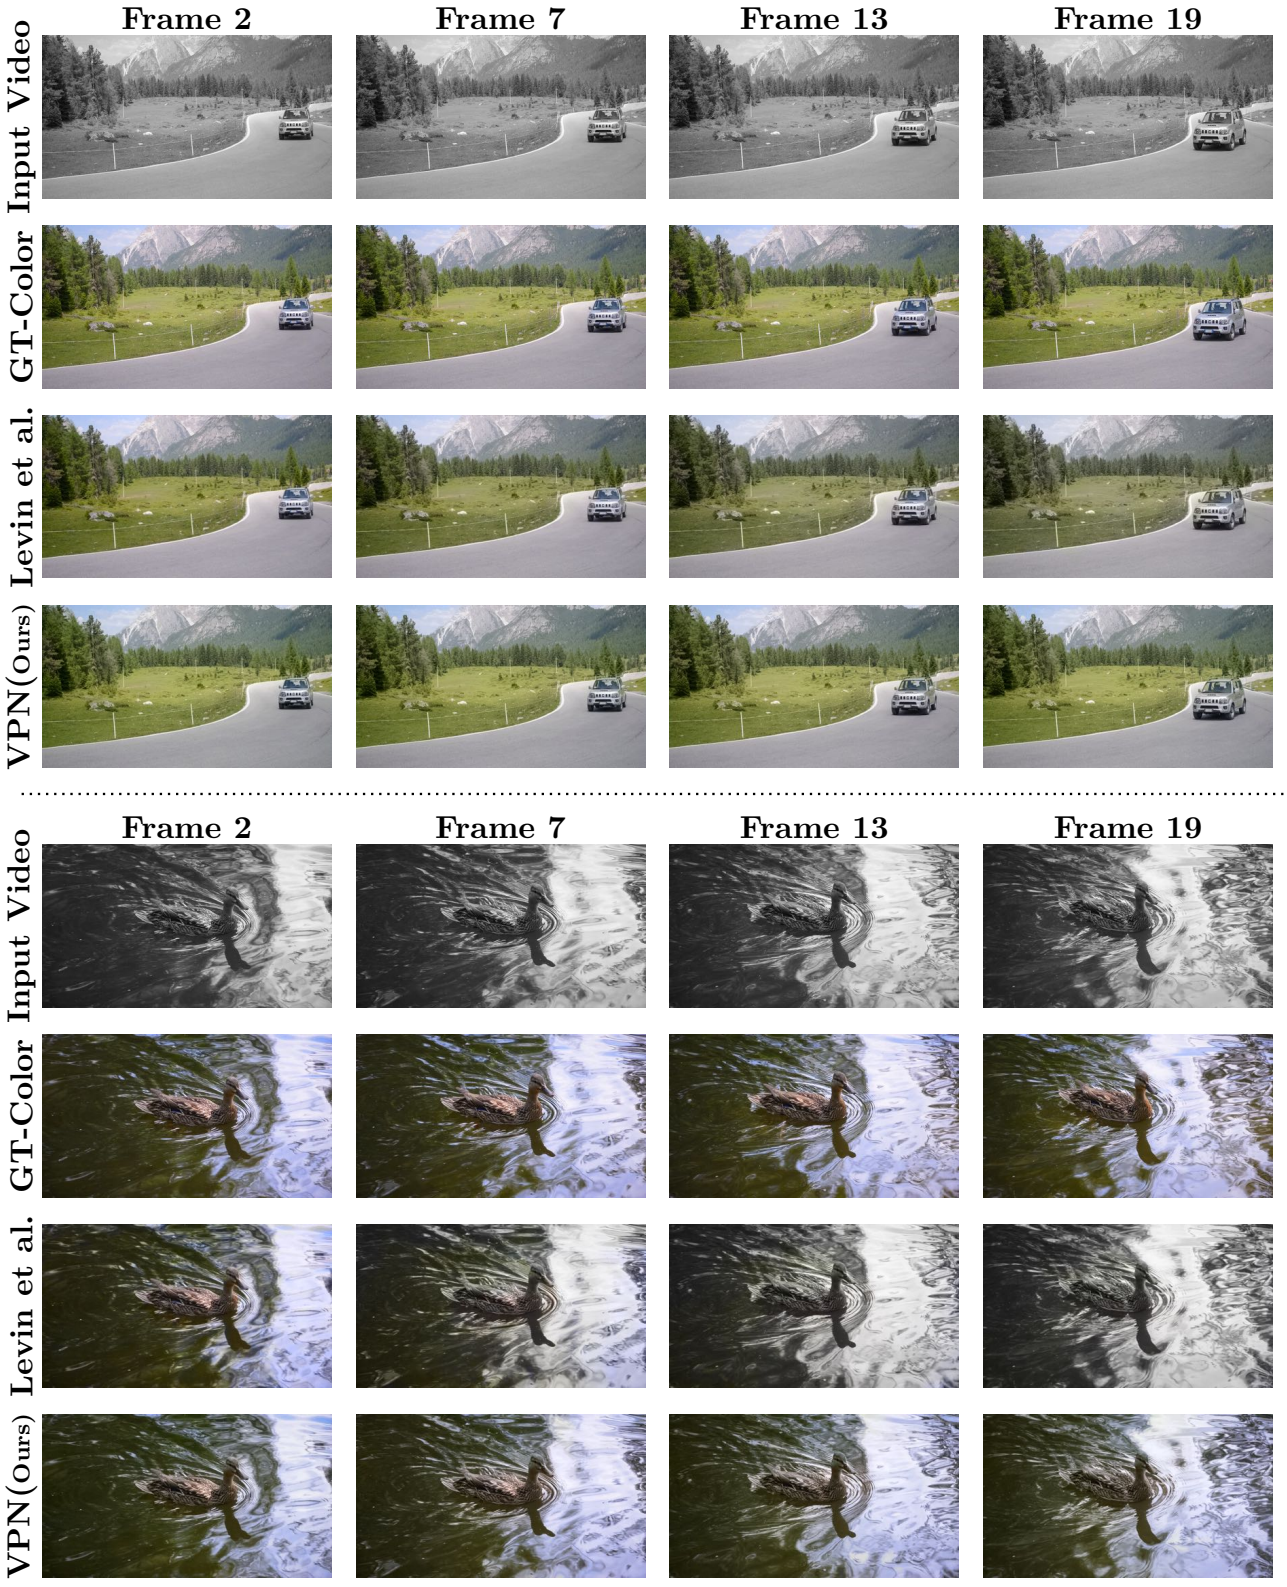

Figure A.5. **Video Color Propagation.** Input grayscale video frames and corresponding ground-truth (GT) color images together with color predictions of Levin et al. [46] and VPN-Stage1 models.